

Programming Stigmergic Coordination with the TOTA Middleware

The Authors

Authors' Affiliations

name@mysite.mycountry

Abstract

Stigmergic coordination has received a growing attention in the past few years. In fact, by decoupling interacting agents via the mediation of an active environment, stigmergy promotes the definition of robust and adaptive multiagent systems. However, beside a large amount of scientific studies, the problem of defining usable and general-purpose tools to program stigmergy-coordinated MAS is still open. In this context, this paper shows how the TOTA middleware can be effectively exploited to support a variety of stigmergy-based coordination activities. The key idea in TOTA is to rely on a simple API for injecting tuple-based information in a network, have it propagate and/or evaporate accordingly to application-specific policies, and have it locally sensed by application agents. Application examples are presented to show that TOTA can promote a simple programming of a variety of different types of stigmergic interactions, suitable for a variety of target operational environments.

Keywords: Stigmergy, Field-based Coordination, Complex MAS, Middleware.

1. Introduction

Stigmergy is being more and more recognized as a powerful approach to coordinate activities in complex multiagent systems [Par97]. Getting its inspiration from nature, stigmergy has proved useful to enforce robust and self-adaptive situated behaviors in a variety of scenarios, ranging from P2P systems [BabMM02], robot swarms [Par04] and sensor networks [Est02].

In general, stigmergy refers to all those kinds of indirect interactions occurring among situated agents that, by affecting and sensing the properties of a shared environment, reciprocally affect each others' behavior [Gra59]. Such indirection makes stigmergic approaches intrinsically suited to large-scale and open multiagent systems, promoting self-organization. Also, by considering the possibility for an active environment to host internal processes affecting its own properties, stigmergy may support powerful forms of context-aware coordination, suitable to tackle the dynamics of modern scenarios and promoting self-adaptation. So far, the most widely used stigmergic mechanisms in multiagent systems include pheromone-based (getting inspiration

from ant foraging, and relying on agents depositing markers that the environment can diffuse and evaporate) [Par97] and field-based ones (getting inspiration from physical force fields and relying on agents' being associated to a sort of "aura" that propagates in the environment) [MamZL04].

Despite the large number of potential applications, practical and usable tools for programming and supporting stigmergic coordinated applications are still missing. Besides scientific simulation studies, where such an issue is not of key relevance, those systems and applications deployed so far that exploit some forms of stigmergic coordination have always adopted specific ad-hoc programming and infrastructural solutions, without attempting at generalizing. However, as the interest in developing and deploying stigmergy coordinated systems increases, the need for general purpose and usable tools will soon become compulsory.

The contribution of this paper is to show how the TOTA ("Tuples On The Air") middleware can represent an effective answer to the above issue. The key ideas in TOTA are to:

- provide basic support for storing, propagation, and maintenance of distributed tuple-based data structures in dynamic networked environments;
- enable a simple application-level definition of tuples, of their propagation rules, and of their maintenance rules (e.g., evaporation rules);
- make available to agents a simple API for injecting tuples in a network (which then propagate and/or "evaporate" accordingly to the defined policies), and have them locally sensed by agents.

In this way – as shown via several application examples – TOTA enable developers and programmers to easily configure any specific type of stigmergic coordination – whether pheromone-based or field-based – and to have their deployment and execution properly supported.

This paper is organized as follows. Section 2 discussed related work in the area. Section 3 introduces the TOTA middleware. Section 4 goes into details about the TOTA API and its programming model. Section 5 sketches several application examples. Section 6 concluded and discusses open issues.

2. Related Work

In the past few years, a growing amount of research activities have been devoted to the study of stigmergic approaches – whether pheromone-based [BonDT99] or field-based [Nag02] – for the coordination of complex multiagent systems.

Beside studies of a more scientific nature exploiting either analytical or simulation approaches [BabMM02, BonDT99], several examples of deployed systems – in a variety of different scenarios and serving different applications – have been reported, e.g.: navigation in P2P networks [TolM04], motion coordination in robot swarms [SveK04, McL04, Par04], data gathering in sensor networks [Est02]; context-aware coordination in pervasive computing environments [MamLZ04], just to mention a few examples. However, despite the recognition of the fact that stigmergic coordination may have wide applicability for a variety of distributed computing scenarios, no proposals so far explicitly focus on the identification and definition of proper general-purpose infrastructures and programming approaches.

The Amorphous Computing project [Abe00], though, focuses on the somewhat close issue of identifying suitable models for programming applications over amorphous networks of “computing particles”. To this end, a simple biologically-inspired programming language based on the propagation and local sensing of simple field-based data structures has been proposed [Nag02]. The project so far has focused on very simple immobile particles, modelled as finite-state machine with a limited number of states - not much different from cellular automata cells. Thus, the effectiveness of the associated programming language to deal with systems of more complex, mobile, and situated particles (as agents in a situated MAS will be) is limited, calling for notable extensions [NagM04]. Also, the focus is on field-based data structures, disregarding issues of e.g., pheromone evaporation.

Shifting to a totally different research area, recent researches on middleware infrastructures for pervasive and mobile computing are proposing a variety of novel coordination abstractions that get somewhat closer to stigmergic coordination than traditional message-passing and event-based middleware do.

In the Smart Messages systems [Bor04], communications between agents/processes occur via sorts of “active” messages that can include code to be executed at each hop in the network path, so to dynamically modify the content of the message itself and/or its routing strategy. In other words, communication between agents involves the active mediation of an active environment, whose activities (encoded in active messages) can be properly configured.

Although, it is potentially possible to program active messages to make them fully assimilated to pheromones and fields, this task requires notable programming efforts.

The L2imbo middleware [Dav01] exploits distributed tuple spaces augmented with internal processes (Bridging Agents) that can move tuples around in the network from one space to another, and can also dynamically change their content. Also in this case, interactions are mediated by an active environment and, by properly defining tuples and bridging agents, one could think at reproducing forms of stigmergic coordination. However, also in this case, the required programming efforts would be notable.

Similar considerations apply to a variety of recently proposed coordination middleware relying on active/reactive components for dynamic data update and propagation, e.g. MARS [Cab03], LIME [PicM04].

The TOTA middleware, described in the following, is explicitly conceived to meet the needs of stigmergic approaches to coordination, thus making it simple to define, program, and access both pheromones and filed, while preserving at the same time generality.

3. The Tuples on the Air Approach

The TOTA middleware gathers concepts from both tuple space coordination architectures [Cab03, PicM04] and event-based ones [CarRW01], but extends them so as to define a simple and flexible coordination mechanisms for general-purpose stigmergic coordination.

3.1 General Overview

In TOTA, we propose relying on distributed tuples for both representing contextual information and enabling stigmergic interactions among distributed agents.

Unlike traditional shared data space models for distributed computing, TOTA tuples are not associated to a specific node (or to a specific data space) of the network. Instead, tuples are “injected” in the network and can autonomously propagate, diffuse, and evolve in the network accordingly to specified patterns. Thus, TOTA tuples form a sort of spatially distributed data structure able to express properties of the network environment that can be used to acquire contextual information about the environment itself, and well as the mechanisms via which stigmergic interactions can take place.

To support this idea, TOTA consider the presence of a peer-to-peer network of nodes, each node running a local version of the TOTA middleware (Figure 1). Each of these TOTA nodes holds references to a limited set of neighbor nodes. The structure of the network, as determined by the neighborhood relations, may be highly

dynamic (due to node mobility, ephemeral nodes, or faults in nodes). TOTA assumes networking capability for recognizing connection and disconnection events.




Figure 1: The TOTA scenario: a peer-to-peer network of nodes, each hosting locally running the TOTA middleware and providing for tuple storing and propagation across the network. At the application level, agents perceive living in an environment in which they can indirectly interact by propagating and locally sensing tuples.

The specific nature of the network scenario determines how each node can find its neighbors, and the overall logical structure of the TOTA network. In MANETs, sensor networks, and robot swarms scenarios, TOTA neighbors nodes are identified within the range of their wireless connection. In wired networks, TOTA neighbors can be determined by some logical structure of the network or by some sort of overlay or social structure. More in general, independently of the network scenario, whenever the nodes of a network can be geographically localized, the structure of the TOTA network can also be based on such spatial information (for the sake of space limitations, we only mention the fact that TOTA can be itself used to enforce geographical self-localization of nodes in a network. The interested reader can refer to [Mam04] for details).

Upon the distributed space identified by the dynamic network of TOTA nodes, which can be considered as the environment in which agents situate, agents can execute on these nodes and exploit the API provided by TOTA to indirectly interact with each other, mostly disregarding network details, and simply injecting tuples from their local position in space and locally sensing propagated tuples.

3.2 TOTA Tuples

Unlike traditional “tuples” which simply have data content, a TOTA tuple is defined in terms of content, a propagation rule, and a maintenance rule:

$$T=(C,P,M)$$

The content C is an ordered set of typed fields representing the information carried on by the tuple. C can range from a simple integer value up to any structured. In this way, TOTA tuples can embed not only simple pheromones and fields, but more semantically reach properties of the environment.

The propagation rule P determines how the tuple should be distributed and propagated across the network. This includes determining the “scope” of the tuple (i.e. the distance at which such tuple should be propagated and possibly the spatial direction of propagation, if not isomorphic) and how such propagation can be affected by the presence or the absence of other tuples in the system. In addition, the propagation rules can determine how tuple’s content should change while it is propagated. In fact, tuples are not necessarily distributed replicas, but can be effectively used to build a distributed overlay data structure expressing some kind of distributed properties of the network environment, e.g., pheromones and fields. For instance, to propagate a field expressing the network distance from the source, a propagation rule can simply propagate the tuples across all the network by increasing a numerical value at each network hop (Figure 2).

The maintenance rule M determines how a tuple’s distributed structure should react to events occurring in the environment. These types of event can be simple time alarms, which are nevertheless of fundamental importance in pheromone-based coordination to enforce evaporation. Or they can be event associated to changes in the network structure, which is of fundamental importance to preserve a coherent structure of the environmental properties represented by field. To this end, the TOTA middleware supports tuples propagation actively and adaptively: by constantly monitoring local events, the network local topology and the income of new tuples, the middleware automatically re-shape tuples and their distributed structure whenever appropriate w.r.t. the maintenance rule. For instance, when new nodes get in touch with a network, TOTA automatically checks the propagation rules of the already stored tuples and eventually propagates the tuples to the new nodes. With reference to the “field” tuple example (Figure 2), the maintenance rule may require that the “field” continuously reflect the actual position of the source agent, whenever it moves. With regard to pheromones, the maintenance rule may require that a tuple is periodically locally update, to enforce evaporation, and that it is periodically propagated, to enforce slow chemical diffusion.

From the agents’ viewpoint point of view, executing and interacting basically reduces to inject tuples, perceive local tuples (as well as “tuple gradients”, by

accessing tuples in neighbor nodes) and local events, and act accordingly to some application-specific policy. Developers, by their side, are charged with the duty of identifying proper content, propagation, and maintenance rules for their application tuples (as, e.g., in Figure 3), and of properly coding them in tuple classes, as described in Section 4.

It is worth mentioning that TOTA, while conceived for stigmergic coordination, also subsumes more traditional forms of indirect interactions. Tuples with a null propagation rule, $T=(C, null, M)$ are simply local, as in reactive tuple space models [Cab03]. Tuples with also a null maintenance rule $T=(C, null, null)$, promote traditional, non-reactive, tuple-based coordination model.

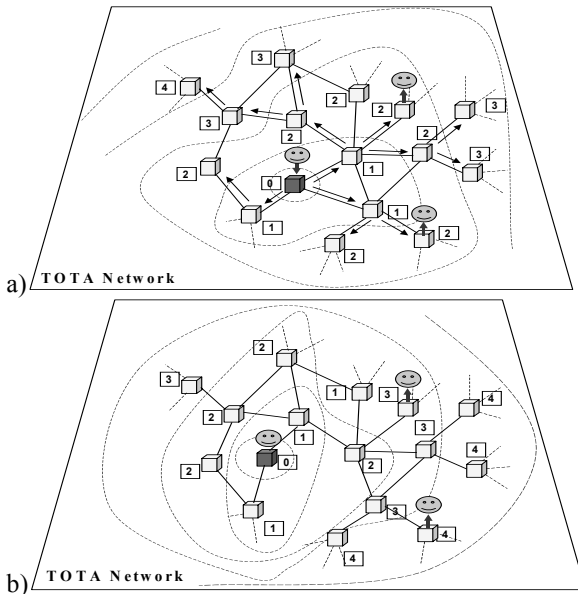


Figure 2: (a) An agent inject a “field” tuple increasing an integer value in its content as it propagates. Other agents can locally perceive the tuples and discover the distance from the source agent. (b) If the source agent (or its node) moves, the maintenance rule can specify to automatically updated the distributed tuple structure to account for the new situation.

Generic Field Tuple

$C = (\text{field_type_identifier}, \text{distance})$

$P = (\text{propagate everywhere, by incrementing distance by one at each network hop})$

$M = (\text{propagate everywhere, by incrementing distance by one at each network hop})$

Generic Pheromone Tuple

$C = (\text{pheromone_type_identified}, \text{strenght})$

$P = (\text{propagate in neighborhood when triggered})$

$M = (\text{evaporate by diminishing strength periodically; trigger propagation periodically})$

Figure 3: Description of generic “field” and “pheromone” tuples, as to be realized in TOTA.

3.3 Implementation

From an implementation point of view, TOTA is fully developed in Java, thus it is highly portable. TOTA can actually run on any JVM-equipped computer device, and can run both in wired and wireless mode. Basically, the implementation includes a tuple space engine, to be locally accessed by agents via an API interface, an event-based engine to catch system-level and network level events, and a reactive engine to execute the code for tuples propagation and maintenance when appropriate events occur.

We extensively experienced the described implementation of TOTA by using Compaq IPAQs, with Linux Familiar and 802.11b in ad-hoc mode, creating the skeleton of the MANET TOTA network.

In addition to the actual implementation of TOTA, we have also implemented a simulator to analyze TOTA behavior in very large systems. The simulator enables examining TOTA behavior in any network scenarios. In addition, the simulator can execute in simulated nodes the *same* TOTA code of real devices, and enable “mapping” in a simulated scenario real network devices. This allow to test applications on a few real devices, while having them behave as if they were immersed in very large networks.

4 TOTA Programming

Developing applications using the TOTA middleware basically implies knowing: (i) what are the primitive operations available in the API to interact with the environment; (ii) how to specify tuples, their propagation rule, and their maintenance rules; (iii) how to properly exploit the above in agents. This latter point will be the core of section 5.

```
public void inject (TotaTuple tuple);
public Vector read (Tuple template);
public Vector readOneHop (Tuple template);
public Tuple keyrd (Tuple template);
public Vector keyrdOneHop (Tuple template);
public Vector delete (Tuple template);
public void subscribe (Tuple template,
    ReactiveComponent comp, String rct);
public void unsubscribe (Tuple template,
    ReactiveComponent comp);
```

Figure 4: The TOTA API

```
public class ToyAgent implements AgentInterface
{
    private TotaMiddleware tota;
    /* agent body */
    public void start() {
        /* create a tuple and inject it*/
        FooTuple foo = new FooTuple("Hello World!");
```

```

tota.inject(foo);
/* define a template tuple */
FooTempTuple t = new FooTempTuple();
/* read local tuples matching the template */
Vector v = tota.read(t);
/* subscribe to changes in tuples matching t*/
tota.subscribe(t, this, "");
}
/* code of the reaction to the subscription */
public void react(String reaction, String
event) { System.out.println(event); }

```

Figure 5: A ToyAgent exploiting the TOTA API

4.1 TOTA Primitives

TOTA is provided with a simple set of primitive operations to interact with the middleware (see Figure 4). *inject* is used to inject the tuple passed as an argument in the TOTA network. Once injected the tuple starts propagating accordingly to its propagation rule (embedded in the tuple definition), and will be stored in each of the propagation nodes in accord to its maintenance rules. The *read* primitive accesses the local TOTA tuple space and returns a collection of the tuples locally present in the tuple space and matching the template tuple passed as parameter. The *readOneHop* primitive returns a collection of the tuples present in the tuple spaces of the node's one-hop neighborhood and matching the template tuple. In stigmergic coordination, such an operation is necessary to estimate "gradients" or either pheromones and fields. To make such an estimation more efficient, the *keyrd* and *keyrdOneHop* are also provided to access tuples based on their unique *id*. The *delete* primitive extracts from the local middleware all the tuples matching the template and returns them to the invoking agent. In addition, *subscribe* and *unsubscribe* primitives are defined to handle events. These primitives rely on the fact that any event occurring in TOTA (including: arrivals of new tuples, connections and disconnections of peers, system-level events) can be represented as a tuple. Thus: the subscribe primitive associates the execution of a reaction method in the agent in response to the occurrence of events matching the template tuple passed as first parameter. Specifically, when a matching event happens, the middleware invokes on the agent a special *react* method passing as parameters, the reaction string and the matching event. The *unsubscribe* primitives removes matching subscriptions

The simple toy agent in Figure 5 clarifies the possible use of the TOTA primitives..

4.2 Specifying Tuples

Being implemented in Java, TOTA tuples are actually objects: the object state models the tuple content, while the tuple's propagation and maintenance rules has been

encoded by means of specific *propagate* and *react* methods, respectively.

When a tuple is injected in the network, it receives a reference to the local instance of the TOTA middleware, then its code is actually executed (the middleware invokes the tuple's *propagate* method) and if during execution it invokes a middleware "move" method, the tuple is actually sent to all the one-hop neighbors, where it will be executed recursively. During migration, the object state (i.e. tuple content) is properly serialized to be preserved and rebuilt upon the arrival in the new host. The *abstract class TotaTuple* provides the basic class on which to rely to define – via inheritance – tuples to serve specific types of stigmergic coordination (Figure 6).

```

abstract class TotaTuple {
protected TotaInterface tota;
/* the state is the tuple content */
...
/* this method inits the tuple, by giving a
reference to the current TOTA middleware */
public void init(TotaInterface tota) {
this.tota = tota; }
/* this method codes the propagation rule */
public abstract void propagate();
/* this method enables the tuple to react to
specific events, to perform maintenance */
public void react(String reaction, String event)
{ }
}

```

Figure 6: The structure of the TotaTuple class

It is worth noting that a tuple is not thread by its own, it is actually executed by the middleware, that runs the tuple's *init*, *propagate*, and *react* methods when necessary. The point to understand is that when the middleware has finished the execution of the tuple's methods, the tuple (on that node) becomes a passive data structure stored in the middleware local tuple space. A tuple is re-activated whenever events for maintenance occur. It is up to tuples themselves to subscribe to interesting events.

While the proposed tuple models provides for the maximum flexibility, coding a specific tuple class from scratch starting from the class *TotaTuple* may be complex. For this reason, TOTA already provides a library of class hierarchies from which the programmer can inherit to create specific tuples without worrying about most of the low-level intricacies of dealing with tuple propagation and maintenance [Auth05]. In particular the basic template from which to inherit is represented by the *StructureTuple* class. *StructureTuple* has been conceived to structure the *propagate* method of *TotaTuple* into four simple and well-defined sub-methods. Subclassing from *StructureTuple* makes the writing of a propagation policy simpler.

StructureTuple implements the *propagate* method accordingly to the schema depicted in figure 7.

In this schema, the four methods *decideEnter*, *makeSubscriptions*, *changeTupleContent* and *decidePropagate* code the tuple behavior:

- *decideEnter()* is executed to assess whether the tuple can enter a specific node
- *makeSubscriptions()* allows the tuple to subscribe to relevant events to perform maintenance operations
- *changeTupleContent()* allows to change the tuple content to create not-dull data-structures
- *decidePropagate()* is executed to asses whether the tuple has to be further propagated to neighbor nodes or not.

Further detail on this schema can be found in [Auth05].

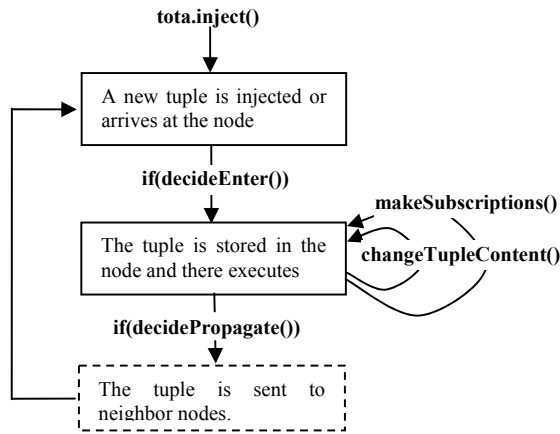


Figure 7: Standard template to create tuples. Tuples are created by overloading *decideEnter*, *makeSubscriptions*, *changeTupleContent* and *decidePropagate* methods.

Tuples can effectively work as stigmergic interaction means only if they are able to maintain their intended distribution despite low-level network topology reconfigurations. These kind of self-maintained tuples are fundamental. In particular, a tuple that maintains a hop-counter expressing hop-distances from the source can be used in a wide range of applications (see 5.1). Given its importance we created this self-maintained *HopTuple* and inserted it in the TOTA library (see figure 8)¹. Inheriting from this tuple, programmers can code other self-maintained tuples effortlessly.

```
public class HopTuple extends StructureTuple
```

¹ For the sake of readability the reported *HopTuple* code has been slightly simplified. Interested readers are encouraged to download the actual code from our website.

```

{
    public int hop = 0;
    /* the tuple subscribes to the incoming of new
    peers (PC REACTION) and to the removal of local
    instances of itself in one-hop neighborhood
    (OUT REACTION) */
    protected void makeSubscriptions() {
        super.makeSubscriptions();
        PresenceTuple pres = new
        PresenceTuple("<peer=*>");
        TsTuple inPres = new
        TsTuple("<op=IN><" + pres.serialize() + ">");
        tota.subscribe(inPres, this, "PC");
        TsTuple tOut = new
        TsTuple("<op=OUT><" + this.serialize() + ">");
        tota.subscribe(tOut, this, "OUT");
    }
    /* the tuple enters there is not a previous
    version of itself or if there is, but with a
    lower-hop */
    protected boolean decideEnter() {
        HopTuple prev = (HopTuple) tota.keyrd(this);
        boolean b1 = prev == null;
        boolean b2 = prev.hop > (this.hop + 1);
        return (b1 || b2);
    }
    protected void changeTupleContent() {
        hop++;
    }
    protected boolean decidePropagate() {
        return true;
    }
    // react method to handle tuple maintenance
    public void react(String react, String event)
    {
        super.react(react, event);
        if (reaction.equalsIgnoreCase("PC"))
        {
            /* a tuple is in a safe-state if it is the
            one originally injected by the agent or if
            it has a neighbor tuple with a lower hop */
            if (safeState() && decidePropagate())
                tota.move(this);
        }
        else if (reaction.equalsIgnoreCase("OUT"))
        {
            if (!safeState())
                tota.delete(this);
            else
                tota.move(this);
        }
    }
}
  
```

Figure 8: The code to realize self-maintained hop-counting tuples.

5 Stigmergy Applications in TOTA

In this section we present two examples to show the expressiveness of the TOTA model to deal with stigmergy in multi-agent-systems. In particular, we will present two exemplary applications one using *gradients* and another using *pheromones*.

5.1 Flocking with Fields

OVERVIEW: The goal of this application is to let a

group of agents to coordinate their movements to maintain a specific distance from each other while moving. To achieve this coordinated behavior, we take inspiration from a well-known example in swarm-intelligence [BonDT99]. Flocks of birds stay together, coordinate turns, and avoid each other, by following a very simple swarm algorithm [Par97]. Their coordinated behavior can be explained by assuming that each bird tries to maintain a specified separation from the nearest birds and to match nearby birds' velocity.

To implement such behavior with TOTA, we can have that each agent generates a tuple *FlockingTuple* whose value assumes the minimal value at specific distance from the source, distance expressing the intended spatial separation between agents (these are network distances, measured in terms of network hops). The final shape of this field approaches the function depicted in figure 9-a. *FlockingTuples* are always updated to reflect peers' movements. To coordinate movements, peers have simply to locally perceive the generated tuples and follow them downhill. The result is a globally coordinated movement in which peers maintain an almost regular grid formation see figure 9-b.




Figure 9: flocking overview. (a) Ideal shape of the flocking tuple. (b) when all the agents follow other agents' tuples they collapse in a regular grid formation.

TUPLES: To code a Flocking tuple, one has to inherit from *HopTuple* and overload the method *changeTupleContent* so as to shape the Flocking tuple value on the basis of the hop-distance, accordingly to figure 9-a. Since we inherit from *HopTuple* and the value is build upon the (self-maintained) hop distance, then also the flocking shape self-maintains.

```
public class FlockingTuple extends HopTuple {
    private int RANGE = 3;
    public FlockingTuple(int RANGE) {
        this.RANGE = RANGE;
    }
    public int value = RANGE;
    protected void changeTupleContent() {
        super.changeTupleContent();
        if(hop <= RANGE)
            value --;
        else
            value ++;
    }
}
```

Figure 10: Flocking tuple

AGENTS: Flocking agents are really simple. They inject flocking tuples then they follow flocking tuple downhill.

```
public class FlockingAgent extends Thread
implements AgentInterface {
    private TotaMiddleware tota;
    public void run() {
        /* inject the flocking tuple to participate the flock */
        FlockingTuple ft = new FlockingTuple ();
        ft.setContent(peer.toString());
        tota.inject(ft);
        while(true) {
            // read other agents' flocking tuples
            FlockingTuple query = new FlockingTuple();
            Vector v = tota.read(query);
            /* evaluate the gradients and select the peer to
            which the gradient goes downhill */
            GenPoint destination = getDestination(v);
            // move downhill following the meeting tuple
            peer.move(destination);
        }
    }
}
```

Figure 11: Agent example: FlockingAgent

5.2 Routing with Pheromones

OVERVIEW: The goal of this application is to realize a routing mechanism for agents in ad-hoc networks based on pheromones. Agents move in an ad-hoc network spreading pheromone trails during their movements. Such pheromone trails can be used to route messages to agents. Visually, we can think the pheromone trail an agent spread as a long tail. Messages can be routed to the agent by following its tail. More in detail, other than the tuples to create pheromones, we created tuples that once injected can follow pheromones up to their source. An agent X willing to send a message to agent Y can wrap the message into one of this follow-pheromone tuples and then inject the tuple.

TUPLES: This application requires two kind of tuples: Pheromone tuples and FollowPheromone tuples.

To code a Pheromone tuple, one has to inherit from *StructureTuple* and realize a maintenance mechanism to let the tuple evaporate after some time. In stark contrast with the field-inspired tuples (*HopTuple* and *FlockingTuple*), Pheromones are local tuples. They spread just one-hop away from the source and they have to remain in that place, before evaporating, even after the agent has moved away. For this reason they do not self-maintain as the field-tuples (see figure 12). Pheromone tuple has been inserted into the TOTA library to be possibly further customized – via inheritance – to create other tuples suited for specific applications.

To code a FollowPheromone tuple, one has to inherit from *StructureTuple* and customize the *decideEnter* method so as the tuple enters only in the nodes having a increasing value of the pheromone (see figure 13).

```

public class Pheromone extends StructureTuple
{
    public int value, VAL, DEC, EVAP;
    public Pheromone(int VAL, int DEC, int EVAP) {
        this.VAL = VAL;
        this.value = VAL;
        this.DEC = DEC; // space decay
        this.EVAP = EVAP; // evaporation
    }
    public void makeSubscriptions() {
        SensorTuple st = new
        SensorTuple("<sensor=clock><value=*>");
        tota.subscribe(st, this, "TIME");
    }
    public boolean decidePropagate() {
        if(value == VALUE)
            return true; // propagate only once
        return false;
    }
    public void changeTupleContent() {
        value = value - DEC;
    }
    public void react(String reaction, String event)
    {
        if(reaction.equalsIgnoreCase("TIME")) {
            value = value -EVAP;
            if(value <= 0) {
                tota.delete(this);
                return;
            }
        }
    }
}

```

Figure 12: Pheromone tuple

```

public class FollowPheromone extends
StructureTuple {
    public int oldVal = 9999;
    Pheromone trail;

    public Pheromone (String msg, String to) {
        content = msg;
        trail = new Pheromone ();
        trail.setContent(to);
    }
    public boolean decideEnter() {
        super.decideEnter();
        int val = getPheromoneValue();
        if(val > oldVal) {
            oldVal = val;
            return true;
        }
        else
            return false;
    }
}

```

Figure 13: Follow pheromone tuple

AGENTS: These kind of pheromone agents are rather simple. They wander randomly spreading a pheromone with an ever increasing value. Then they have a method to send messages to other agents.

```

public class PheromoneAgent extends Thread
implements AgentInterface {
    private TotaMiddleware tota;
    public void run() {
        int val = 10;
        while(true) {

```

```

// move randomly
peer.move(Math.random());
// while spread pheromone
Pheromone p = new Pheromone(val, 5, 1);
p.setContent(peer.toString());
peer.inject(p);
}
}
public void send(String msg, String to) {
    FollowPheromone fp = new
    FollowPheromone(msg, to);
    tota.inject(fp);
}
}

```

Figure 14: Agent example: the agent moves randomly spreading pheromones. Moreover, it has a method to send a message, wrapped in a FollowPheromone tuples to another agent.

6 Conclusions and Future Works

Stigmergy is getting more and more recognized as a relevant approach for supporting the definition of robust and self-adaptive multiagent systems. Still, little has been done so far to leverage the practically exploitation and deployment of stigmergy-coordination multiagent systems. The TOTA middleware, by making available to simple API with which to program, in an effective way, a number of diverse stigmergic coordination patterns, proposes itself as a general-purpose approach for programming complex multiagent systems.

Despite the potentials of TOTA, several issues still need to be faced to increase its usability. First, security issues, disregarded in most researches in stigmergic coordination can no longer be ignored for systems which are to be deployed in an open and possibly hostile environment. However, what stigmergy implies in terms of security and privacy is to be fully explored. Second, the lack of an underlying general methodology, enabling engineers to map a specific coordination pattern into the corresponding definition of tuples and of their propagation/evaporation rules, is still to be identified. Nevertheless, this is a general drawback of researches on complex MAS, rather than a specific drawback of the TOTA approach.

References

- [Auth04] The Authors. "Programming Pervasive and Mobile Computing Applications with the TOTA Middleware", *2nd IEEE Conference on Pervasive Computing and Communication*, Orlando (FL), March 2004.
- [Auth05] The Authors. "Self-Maintained Distributed Tuples for Field-based Coordination in Dynamic Networks", *Concurrency and Computation: Practice and Experience*, 2005, to appear.
- [Abe00] H. Abelson, et al., "Amorphous Computing", *Communications of the ACM*, 43(5), May 2000.

- [BabMM02] O. Babaoglu, H. Meling, A. Montresor, "Anthill: a Framework for the Development of Agent-Based Peer-to-Peer Systems", *Proceedings of the 22nd IEEE Conference on Distributed Computing Systems*, Vienna (A), May 2002.
- [BonDT99] E. Bonabeau, M. Dorigo, G. Theraulaz, "Swarm Intelligence", Oxford University Press, 1999.
- [Bor04] C. Borcea, "Spatial Programming Using Smart Messages: Design and Implementation", *24th Int.l Conference on Distributed Computing Systems*, Tokio (J), May 2004.
- [Cab03] G. Cabri, L. Leonardi, M. Mamei, F. Zambonelli, "Location-dependent Services for Mobile Users", *IEEE Transactions on Systems, Man, and Cybernetics*, 33(6):667-681, Nov. 2003
- [CarRW01] A. Carzaniga, D. Rosenblum, A. Wolf, "Design and Evaluation of a Wide-Area Event Notification Service", *ACM Transaction on Computer System*, 19(3):332-383, 2001.
- [Dav01] N. Davies, et al, "L2imbo: A distributed systems platform for mobile computing", *ACM Mobile Networks and Applications*, 3(2):143-156, 2001.
- [Est02] D. Estrin, D. Culler, K. Pister, G. Sukjatme, "Connecting the Physical World with Pervasive Networks", *IEEE Pervasive Computing*, 1(1):59-69, 2002.
- [Gra59] P.-P. Grassé. La Reconstruction du Nid et les Coordinations Inter-Individuelles chez *Bellicositermes Natalensis* et *Cubitermes* sp. La théorie de la Stigmergie: Essai d'Interprétation du Comportement des Termites Constructeurs. *Insectes Sociaux*, 6:41-84, 1959.
- [MamZL04] M. Mamei, F. Zambonelli, L. Leonardi, "Co-Fields: A Physically Inspired Approach to Distributed Motion Coordination", *IEEE Pervasive Computing*, 4(2):52-61, 2004.
- [MasCE01] C. Mascolo, L. Capra, W. Emmerich, "An XML based Middleware for Peer-to-Peer Computing", *IEEE Intl. Conference of Peer-to-Peer Computing*, 2001.
- [McL04] J. McLurkin, J. Smith, "Distributed Algorithms for Dispersion in Indoor Environments using a Swarm of Autonomous Mobile Robots", *Proceedings of the 7th International Symposium on Distributed Autonomous Robotic Systems*, Toulouse (F), 2004.
- [Nag02] R. Nagpal, "Programmable Self-Assembly Using Biologically-Inspired Multi-agent Control", *1st Int.l Conference on Autonomous Agents and Multi-agent Systems*, Bologna (I), July 2002.
- [NagM04] R. Nagpal, M. Mamei, "Engineering Amorphous Computing Systems", in *Methodologies and Software Engineering for Agent Systems: the Handbook of Agent-Oriented Software Engineering*, Kluwer Academic Publishing (New York, NY), 2004.
- [Par97] V. Parunak, "Go to the Ant: Engineering Principles from Natural Agent Systems", *Annals of Operations Research*, 75:69-101, 1997.
- [Par04] V. Parunak, S. Brueckner, J. Sauter, "Digital Pheromones for Coordination of Unmanned Vehicles", *Workshop on Environments for Multi-agent Systems (E4MAS)*, LNAI 3374, Springer Verlag, 2004.
- [PicM04] G. P. Picco, A. L. Murphy, "Using Coordination Middleware for Location-Aware Computing: A Lime Case Study", *Proceedings of the 6th International Conference on Coordination Models and Languages, LNCS No. 2949*, Feb. 2004.
- [SveK04] J. Svennebring, S. Koenig, "Building Terrain Covering Ant Robots: a Feasibility Study", *Autonomous Robots*, 16(3):313-332, May 2004.
- [TolM04] R. Tolksdorf, R. Menezes, "Using Swarm Intelligence in Linda Systems", *Proceedings of the 4th International Workshop on Engineering Societies in the Agents' World*, LNCS No. 3071, 2004.
- [ZamP04] F. Zambonelli, V. Parunak, "Towards a Paradigm Change in Computer Science and Software Engineering", *The Knowledge Engineering Review*, Vol. 18, 2004, to appear.